

Abstract

Packets are scheduled for transmission over a communication link in a network, using a Largest Weighted Delay First (LWDF) scheduling policy. A delay measure W_i , $i = 1, 2, \dots, N$, is computed for each of N packets, each associated with a corresponding one of N data flows and located in a head position in a corresponding one of N data flow queues. The computed delay measures are then weighted using a set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$. The packet having the largest weighted delay W_i/α_i associated therewith is then selected for transmission. In an embodiment configured to meet a quality of service (QoS) requirement specified in terms of a deadline T_i and an allowed deadline violation probability δ_i , e.g., a requirement specified by $P(W_i > T_i) \leq \delta_i$, the weights α_i in the set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$ may be given by $\alpha_i = -T_i / \log \delta_i$. The invention can also be used to meet other types of QoS requirements, including, e.g., requirements based on packet loss probabilities. For example, the QoS guarantee may be defined for a delay measure in the form of queue length Q_i , $i = 1, 2, \dots, N$, and an allowed queue length violation probability δ_i . In such an embodiment, the QoS requirement is specified by $P(Q_i > H_i) \leq \delta_i$, and the weights α_i in the set of positive weights $\alpha_1, \alpha_2, \dots, \alpha_N$ may be given by $\alpha_i = -H_i / \log \delta_i$, where H_i represents an upper bound on the length of the queue.